Application Notes

hp StorageWorks Continuous Access EVA Replication Performance Estimator V3.0

Product Version: 3.0

Fourth Edition (July 2004)

Part Number: AA-RU5MD-TE

This document provides instructions for the use of the replication performance estimator, which provides a way to analyze the effects of distance on applications that use HP StorageWorks Continuous Access Enterprise Virtual Array (EVA)

For the latest version of these Application Notes and other Continuous Access EVA documentation, access the HP storage web site at: http://www.hp.com/country/us/eng/prodserv/storage.html.

Printed in the U.S.A.

Continuous Access EVA Replication Performance Estimator V3.0 Application Notes
Fourth Edition (July 2004)
Part Number: AA-RU5MD-TE

About this Document

This application note covers the following topics:

- The Continuous Access EVA Performance Estimator, page 4
- Using the Continuous Access EVA Performance Estimator, page 4
- Inputs to the Estimator, page 5
- Estimator results, page 10
- Best practices, page 16
- Glossary, page 18

Audience

This document is intended for customers who are using Continuous Access Enterprise Virtual Array (EVA) and need to estimate the effects of distance on applications that use the Continuous Access EVA.

Other Continuous Access EVA Documentation

In addition to this guide, HP provides information related to Continuous Access EVA in the following:

- *HP StorageWorks Continuous Access EVA V1.1 Design Reference Guide*, part number AA-RS2YD-TE
- *HP StorageWorks Continuous Access EVA V1.1 Getting Started Guide,* part number T3031-96201
- *HP StorageWorks Continuous Access User Interface V1.1 Installation Guide,* part number T3031-96103
- *HP StorageWorks Continuous Access EVA V1.1 Operations Guide*, part number AA-RTEHC-TE
- *HP StorageWorks Command View EVA V3.1 Getting Started Guide*, part number AA-RQZBE-TE
- *HP OpenView Storage Operations Manager Installation Overview*, part number T2538-99001
- *HP StorageWorks Enterprise Virtual Array User Guide EVA3000*, part number EK-EVA30-UG. B01
- *HP StorageWorks Enterprise Virtual Array License Instructions*, part number AA-RS2AB-TE
- *HP StorageWorks SAN Design Reference Guide*, part number AA-RMPNH-TE
- *HP StorageWorks Continuous Access EVA and Data Replication Manager SAN Extension Reference Guide*, part number AA-RU5CB-TE
- HP StorageWorks Continuous Access User Interface V1.1 Online Help
- HP StorageWorks Command View EVA V3.1 Online Help

Additional documentation, including white papers and best practices documents, are available on the HP website at http://h18006.www1.hp.com/storage/software.html.

The Continuous Access EVA Performance Estimator

When the Microsoft Excel-based performance is first loaded it appears as shown in Figure 1.

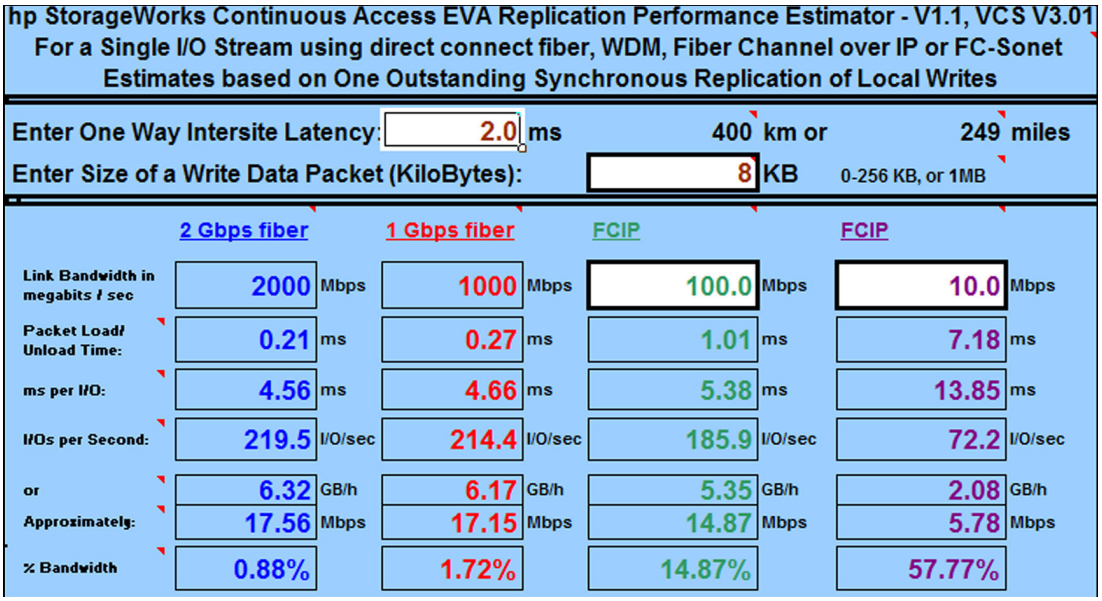

hp StorageWorks Continuous Access EVA Replication Performance Estimator - V1.1, VCS V3.01
For a Single I/O Stream using direct connect fiber, WDM, Fiber Channel over IP or FC-Sonet
Estimates based on One Outstanding Synchronous Replication of Local Writes

Enter One Way Intersite Latency: 2.0 ms 400 km or 249 miles
Enter Size of a Write Data Packet (KiloBytes): 8 KB 0-256 KB, or 1MB

	2 Gbps fiber		1 Gbps fiber		FCIP		FCIP	
Link Bandwidth in megabits / sec	2000	Mbps	1000	Mbps	100.0	Mbps	10.0	Mbps
Packet Load/ Unload Time:	0.21	ms	0.27	ms	1.01	ms	7.18	ms
ms per I/O:	4.56	ms	4.66	ms	5.38	ms	13.85	ms
I/Os per Second:	219.5	I/O/sec	214.4	I/O/sec	185.9	I/O/sec	72.2	I/O/sec
or	6.32	GB/h	6.17	GB/h	5.35	GB/h	2.08	GB/h
Approximately:	17.56	Mbps	17.15	Mbps	14.87	Mbps	5.78	Mbps
% Bandwidth	0.88%		1.72%		14.87%		57.77%	

Figure 1: Performance Estimator

Using the Continuous Access EVA Performance Estimator

The performance estimator provides a way to examine the effects of distance on applications that use Continuous Access EVA. The mathematics and much of the details embedded within the spreadsheet are described in the *HP StorageWorks Continuous Access EVA Design Reference Guide*, which is available from the same website where you acquired the estimator: http://h18006.www1.hp.com/storage/software.html

The estimator is designed to support both direct fiber and any FC-Sonet or Fibre Channel IP-based intersite link technology, such as those shown on the right side of Figure 2, and will accept any bandwidth from 1 Mbps to 1000 Mbps.

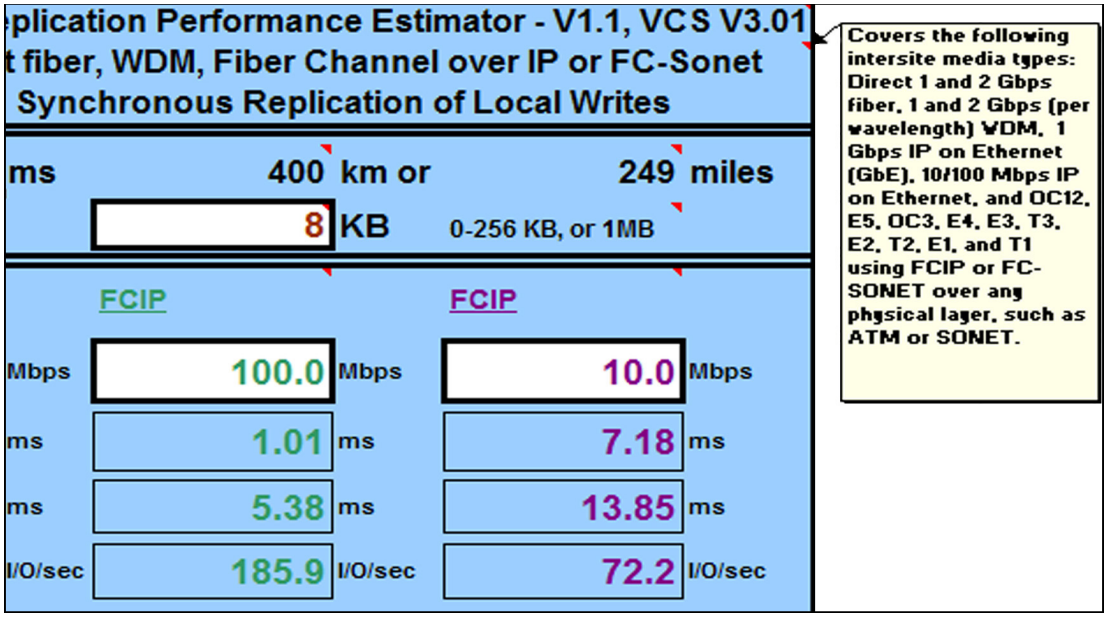

Figure 2: Applicable media types

Table 1 shows the bandwidths of the various intersite links listed in Figure 2.

Table 1: Bandwidths of intersite links

Intersite Link Technology	Intersite Link Bandwidth (Mbps)
2-Gbps Fibre Channel	2000
1-Gbps Fibre Channel	1000
1-GbE (Gigabit Ethernet) IP	1000
OC12 IP	622
E5 IP	565
OC3 IP	155.5
E4 IP	139.3
100-Mbps Ethernet IP	100
T3 IP	44
E3 IP	34.304
10-Mbps Ethernet IP	10
E2 IP	8.4
T2 IP	6.2
E1 IP	2.048
T1 IP	1.54

Inputs to the Estimator

Two inputs are required for an estimate:

- **One-way Intersite Latency.** This is one-half of a network PING or is estimated using the rules of thumb shown in the comment attached to the cell (see Figure 3). Figure 4 and Figure 5 show the latency in kilometers and miles, respectively, of optical cable.
- **Packet Size.** This is the average size of the write that will be replicated by Continuous Access EVA to the array located "One-Way Intersite Latency" away from the source EVA (see Figure 6 and Figure 7). If there are to be multiple sizes of writes, due, for example, to multiple applications, then run the estimate for each different size of write particular to an application. For an example of how to combine multiple size writes see "Performance Impact of Multiple Replication Relationships" in the *HP StorageWorks Continuous Access EVA Design Reference Guide*.

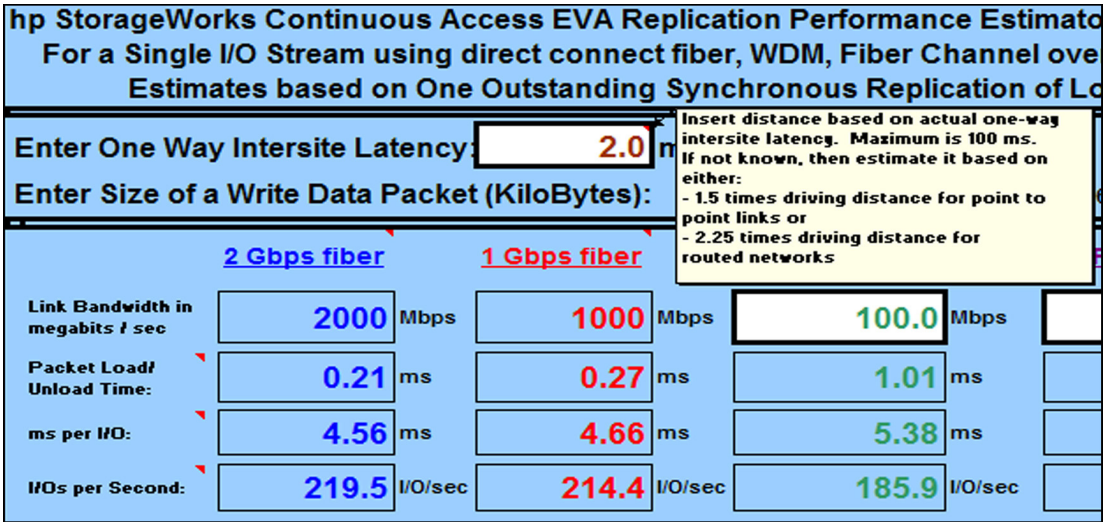

Figure 3: One-way latency

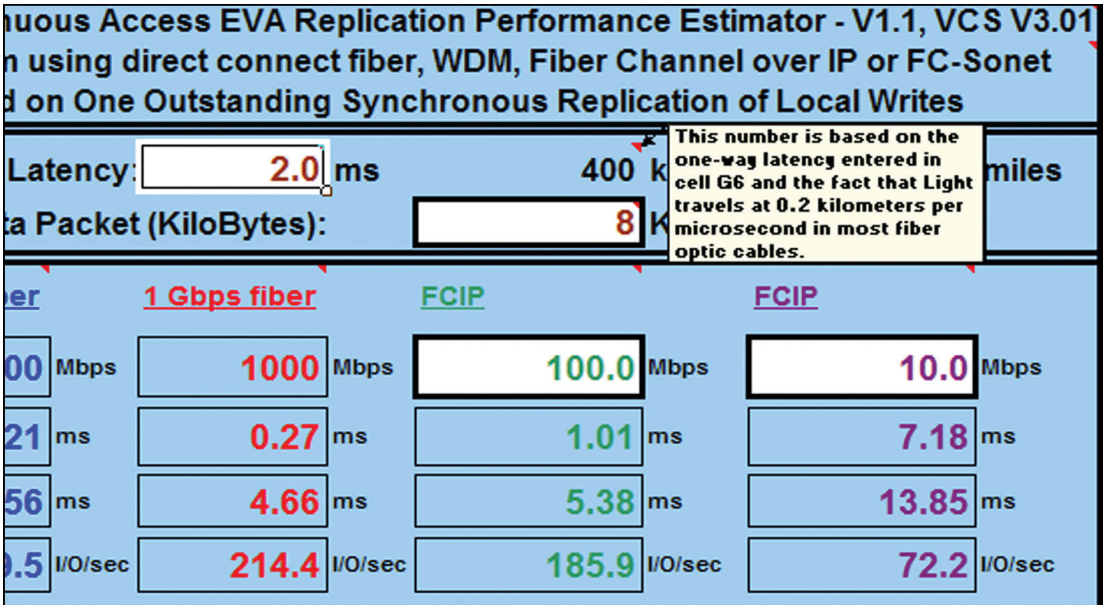

Figure 4: One-way latency in kilometers

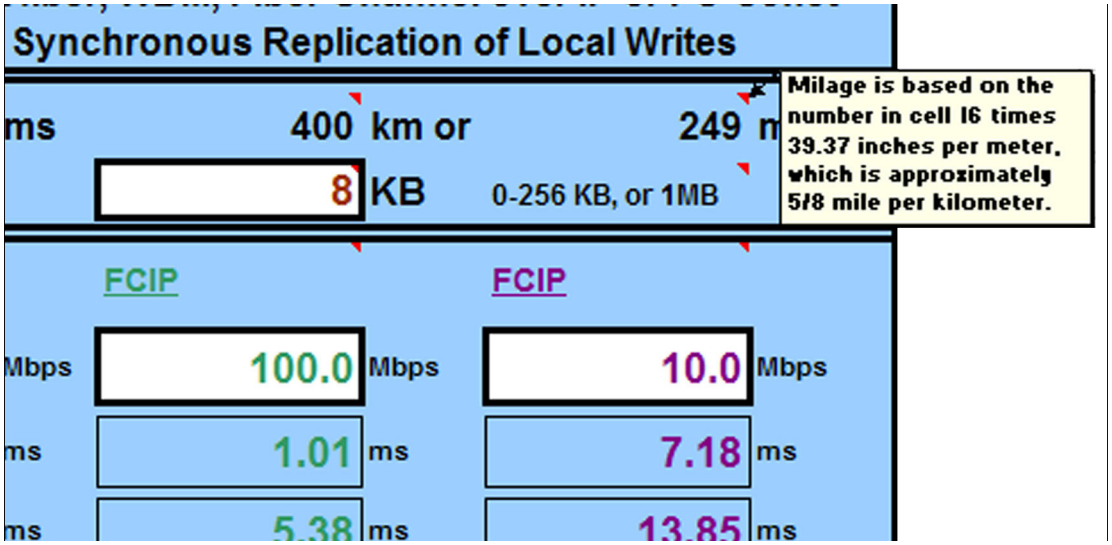

Figure 5: One-way latency in miles

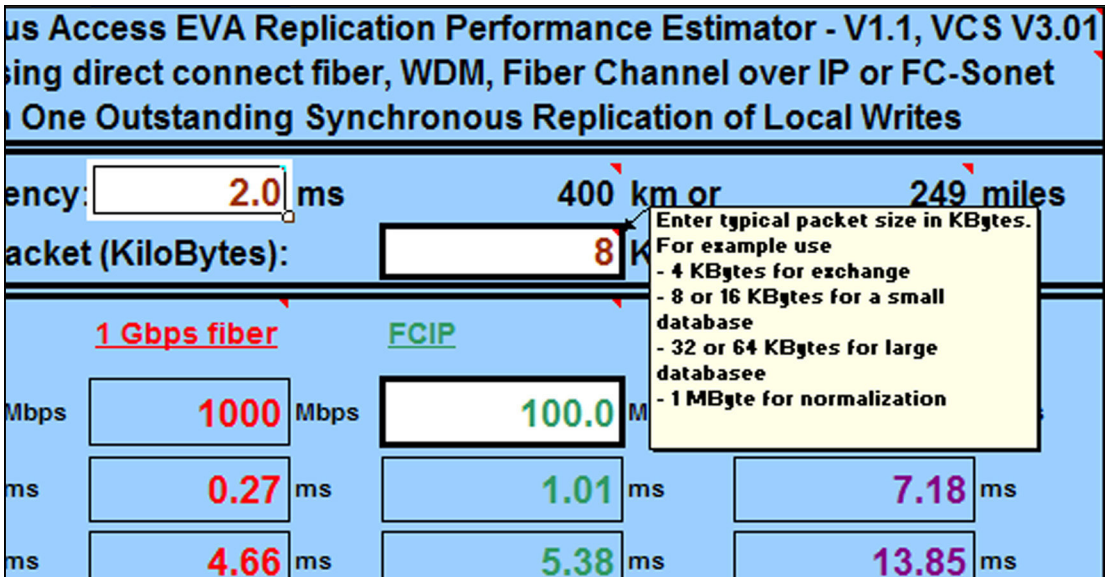

Figure 6: Packet size

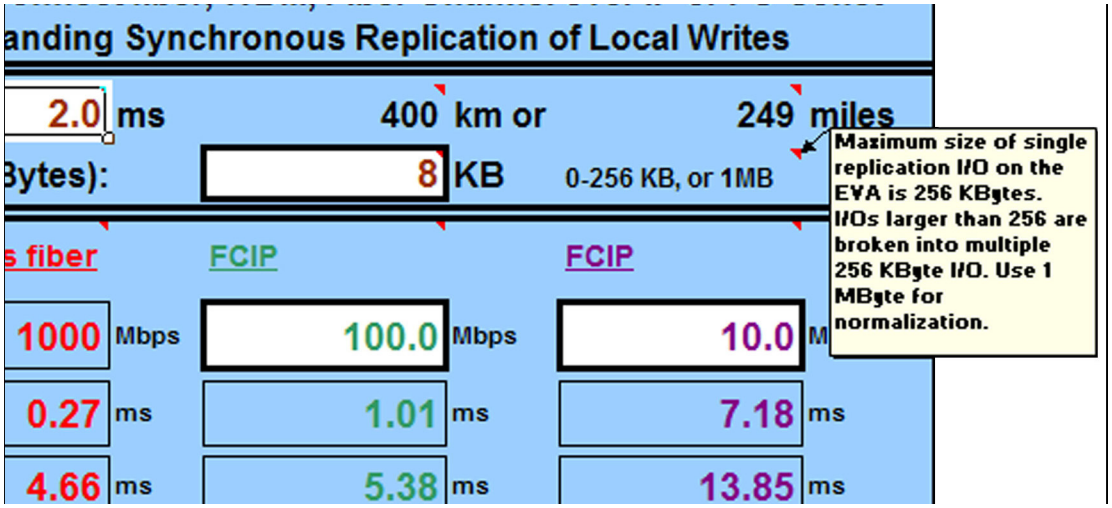

Figure 7: Maximum packet size

Two optional inputs are also available. These are the two columns that default to the FCIP bandwidth values of 100 and 10 Mbps (see Figure 8 and Figure 9). These two values can be changed as needed to show the effects of changes in bandwidth on replication performance. As a starting point, HP recommends using the defaults, then using the values provided in Tables 2 and 3, and then, as the analysis proceeds, adjusting one of the values up or down while leaving the other at the default value for comparison.

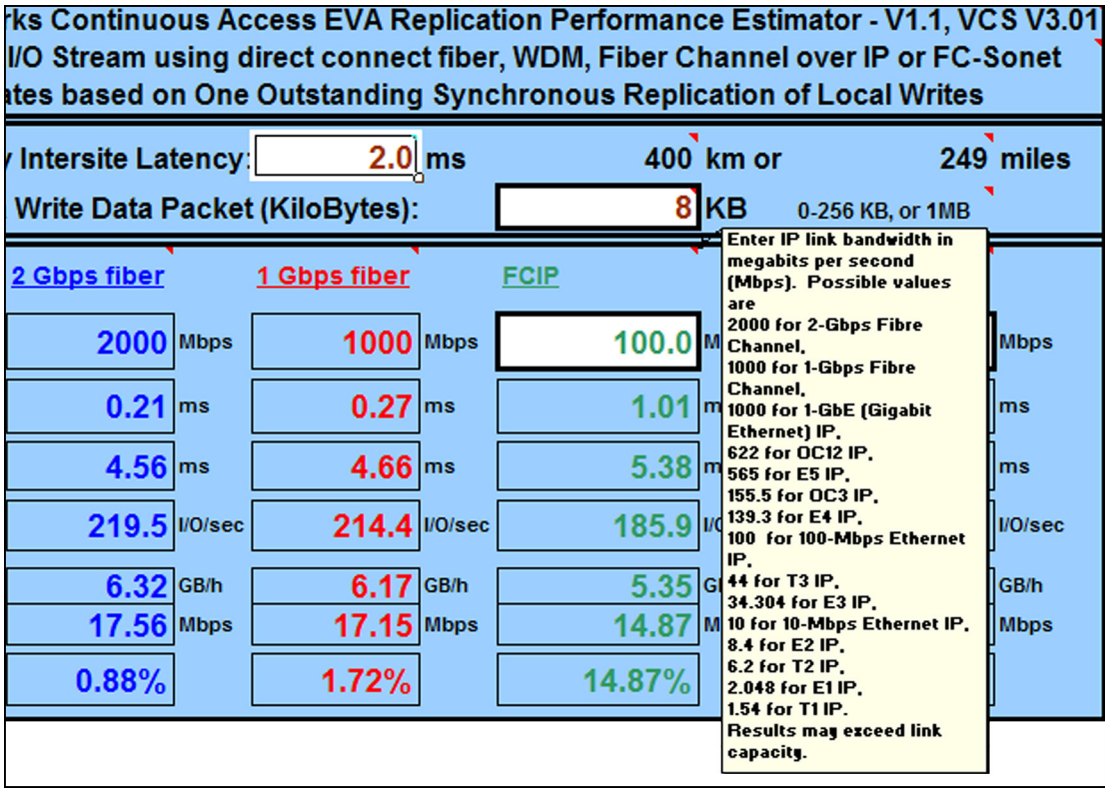

Figure 8: Input IP link bandwidth and its output

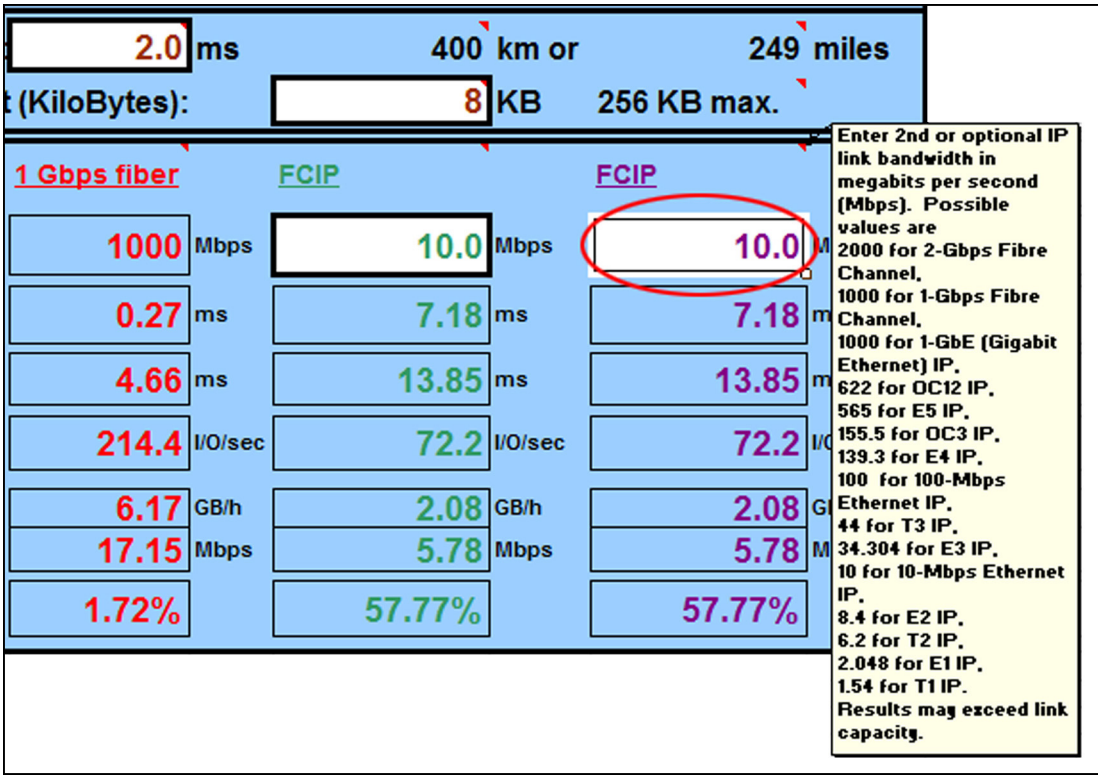

Figure 9: Input optional IP link bandwidth and its output

In the first row of the estimator, the 1- and 2-Gbps fiber bandwidths are fixed. These two columns support the intersite link shown in Figure 10 and Figure 11. Both columns are to be used for any long-distance direct fiber connection using longwave or very long distance GBIC/SFP, of either coarse- or dense-wave division multiplexing (CWDM or DWDM).

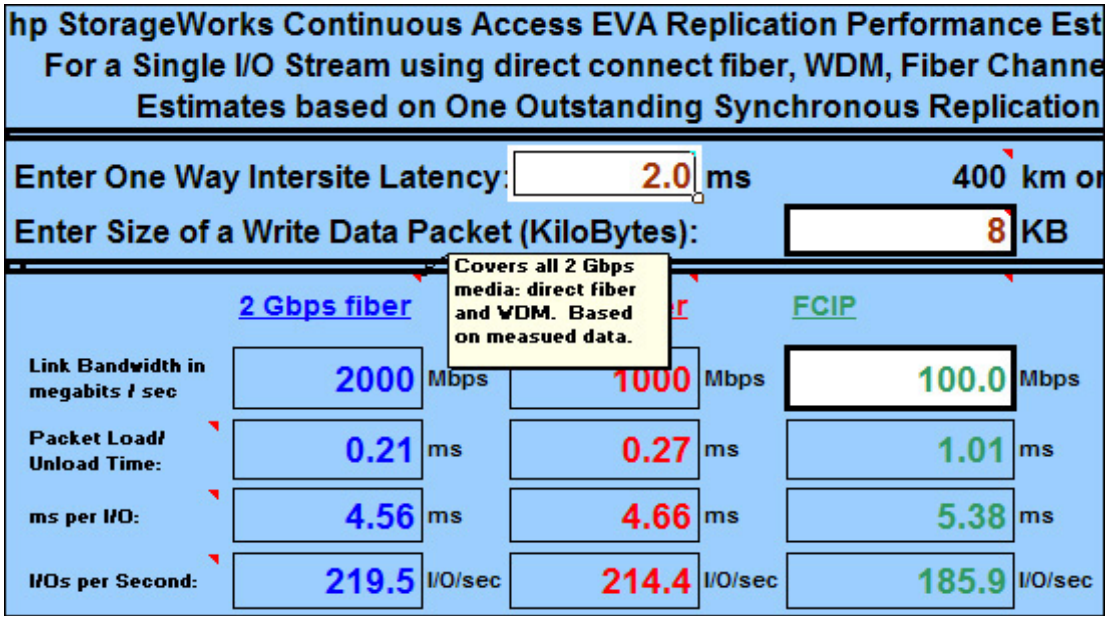

Figure 10: Output column for 2-Gbps fiber

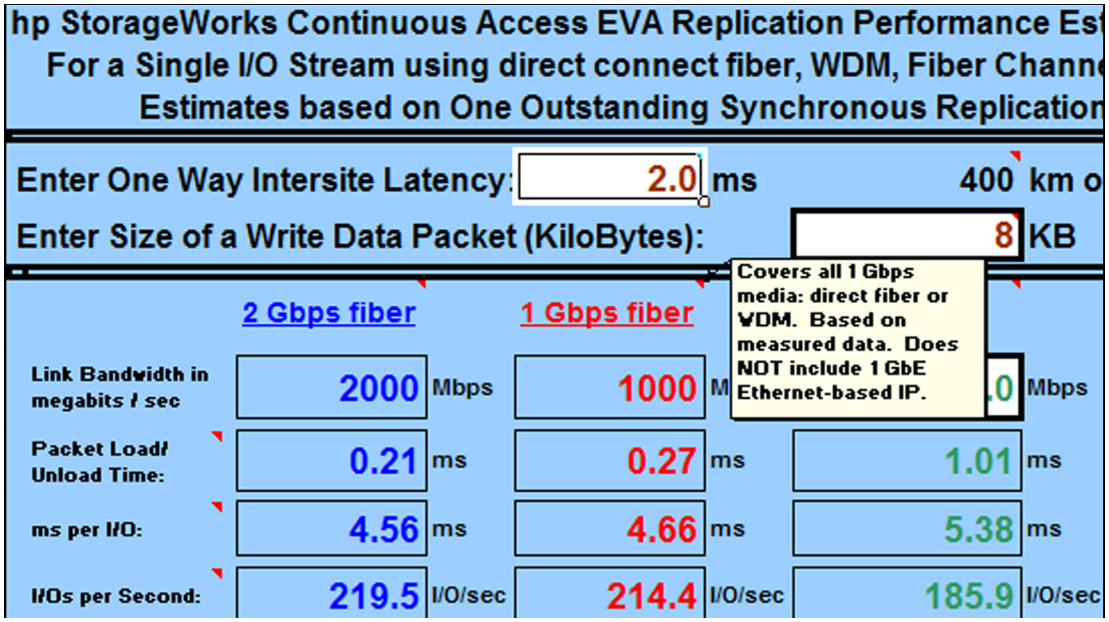

Figure 11: Output column for 1-Gbps fiber

Estimator results

For each set of one-way latency and packet size, the estimator produces six rows of output over the four columns. The six output rows are:

- **Packet Load/Unload Time.** This is the time it takes for a packet to be loaded on to and off of the intersite link (see Figure 12). For the purpose of the estimator, packet load/unload time is based on the bandwidth of the intersite link, which is usually the narrowest part of the solution (from a bandwidth perspective). Figure 13 shows the effect of bandwidth by comparing the length in time of a packet in both T3 and OC3 links.

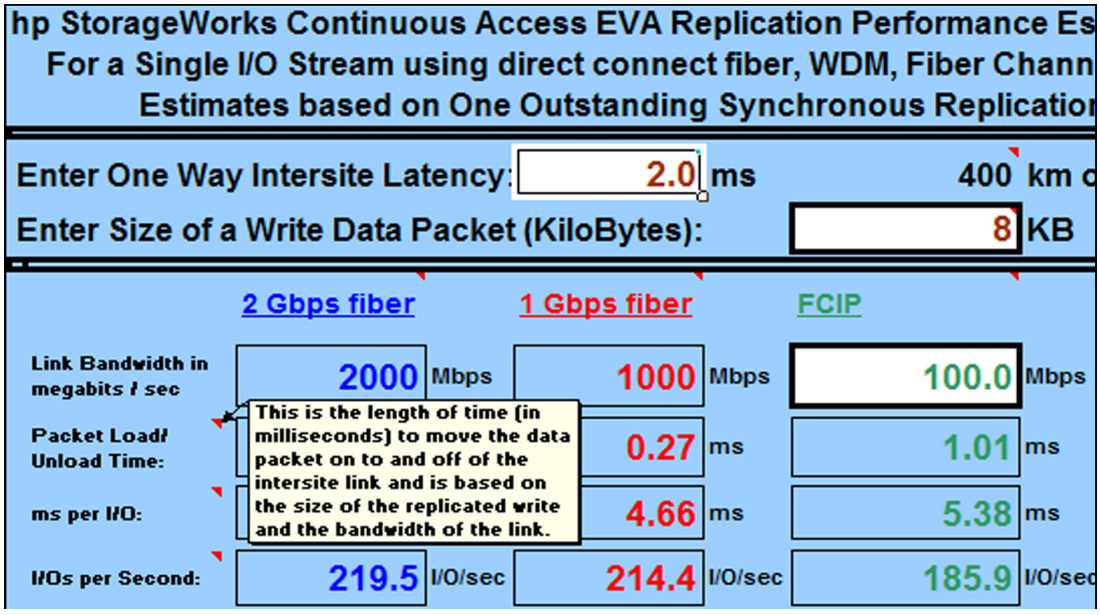

Figure 12: Output row: packet transmit time

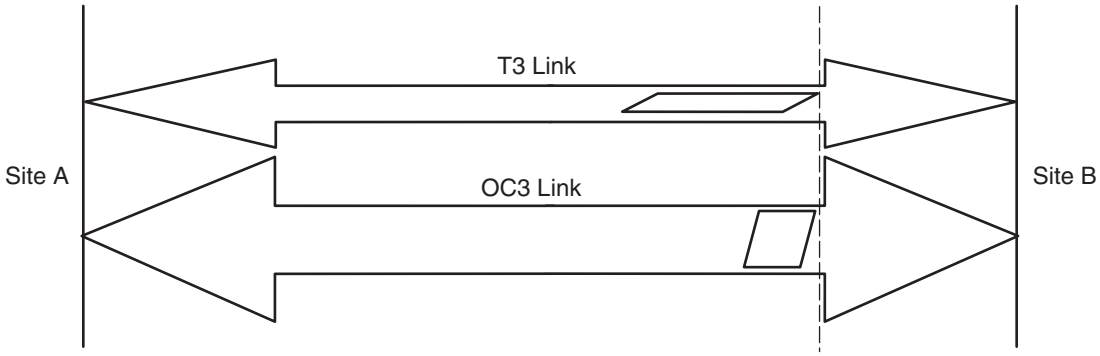

Figure 13: Effect of bandwidth on packet length

- **ms per I/O.** This is the sum of the packet load/unload time, twice the one-way intersite latency, and the replication overhead (see Figure 14). The value returned is the time it takes to complete a single synchronous write for a given bandwidth. In the *HP StorageWorks Continuous Access EVA Design Reference Guide*, the mathematics that calculates this row is shown as: $y = mx + b$, where
  - y is the time for the single I/O.
  - m is the slope of the line that is representative of the bandwidth of the link.
  - x is the size of the transfer.
  - b is the intercept of the slope with the y-axis and is representative of the replication overhead for a single 512-byte write.
- **I/Os per Second.** Inverting the time to complete a single write results in results in the theoretical peak writes per second (see Figure 15). This number can be achieved only if there is no delay between the completion of one write and the start of the next. In a normal application where writes are generated in a pseudo-random fashion, the expected peak would be approximately 70% of the theoretical peak. The average rate typically does not exceed 50% of the theoretical peak. Table 2 shows the I/Os per Second of various SAN technologies.
- **Throughput in GB per hour.** This value is calculated by multiplying the I/O per second by the size of the write and then by 3600 seconds per hour for the given set of links and intersite separation (see Figure 16).
- **Throughput in megabits per second (Mbps).** This value is calculated by dividing the GByte per hour by 3600 seconds per hour and then multiplying by 10 bits per byte to obtain the expected peak serial data flow (see Figure 17). Table 2 shows the throughput of various SAN technologies.
- **Percent Bandwidth.** This value is calculated by dividing the throughput in Mbps by the link bandwidth at the top of the column and then multiplying by 100% (see Figure 18). Table 2 shows the single stream bandwidth of various SAN technologies, and Table 3 shows multiple stream bandwidth of various SAN technologies.

Table 2: I/O Per Second, Throughput, and Percent Bandwidth Used for 2-KB Writes

Technology	Approximate IOPS	Throughput (Mbps) based on 2 KB per I/O	Approximate Single Stream Percent Bandwidth Used
2-Gbps Fibre Channel	417.50	8.35	0.4%
1-Gbps Fibre Channel	405.60	8.11	0.8%
1-GbE (Gigabit Ethernet) IP	403.13	8.06	0.8%
OC3 IP	393.44	7.87	5.1%
E4 IP	391.97	7.84	5.6%
100-Mbps IP	386.01	7.72	7.7%
T3 IP	361.40	7.23	16.1%
E3 IP	348.77	6.98	20.5%
10-Mbps IP	246.08	4.92	49.2%
E1 IP	82.91	1.66	82.9%
T1 IP	64.77	1.3	86.4%

Table 3: Example Summary Data for Multiple Applications Writing Across the Same Link

Technology	32 KB IOPS times 3	Throughput (Mbps) based on 32 KB per I/O	Percent Bandwidth Used	2 KB IOPS times 4	Throughput (Mbps) based on 2 KB per I/O	Percent Bandwidth Used	Percent Total Bandwidth Required
2-Gbps Fibre Channel	938	300	15%	1670	33	2%	17%
1-Gbps Fibre Channel	867	277	28%	1623	32	3%	31%
1-GbE (Gigabit Ethernet) IP	859	274	27%	1613	32	3%	31%
OC3 IP	623	199	129%	1574	31	20%	149%
E4 IP	599	192	138%	1568	31	23%	161%
100-Mbps IP	522	167	167%	1544	31	31%	198%
T3 IP	334	107	238%	1446	29	64%	302.%
E3 IP	276	88	260%	1395	28	82%	342%
10-Mbps IP	98	31	313%	984	20	197%	510%
E1 IP	20	6.6	330%	332	6.6	332%	662%
T1 IP	15	5	330%	259	5.2	345%	676%

Note: There will be a slight difference between the results generated using the estimator and those shown in tables 2 and 3. The main reason for this is that the spreadsheet uses a mathematical model based on the IP link bandwidth to estimate the slope and intercept of the line that is used in the first two rows of calculations. In tables 2 and 3 experimentally derived data is used for each type of link because there is a nonlinear relationship between slope and bandwidth and between y-intercept and bandwidth.

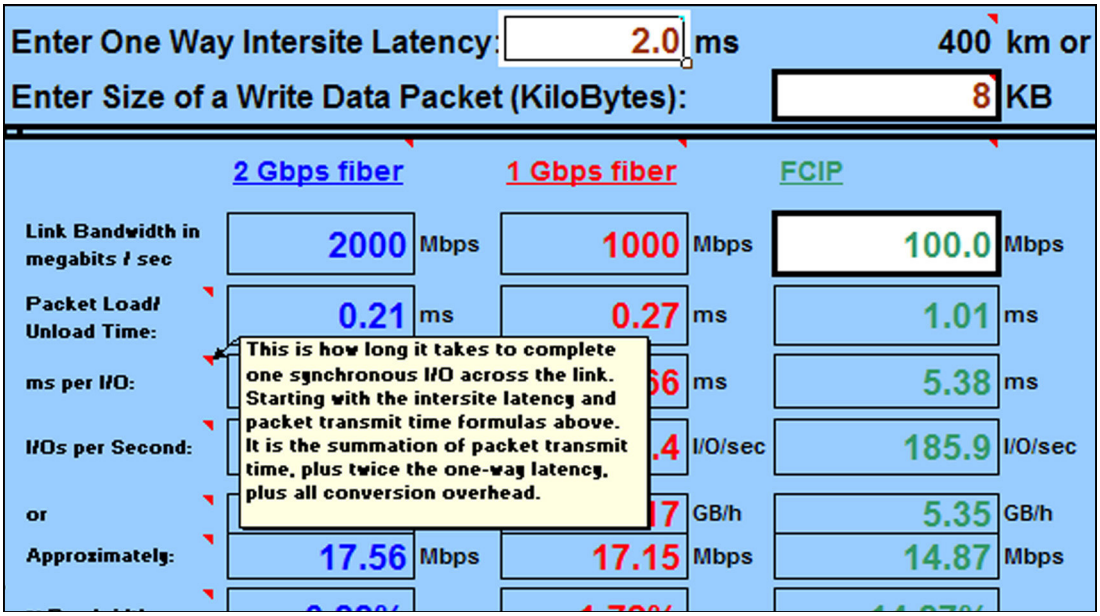

Figure 14: Output row: ms per I/O

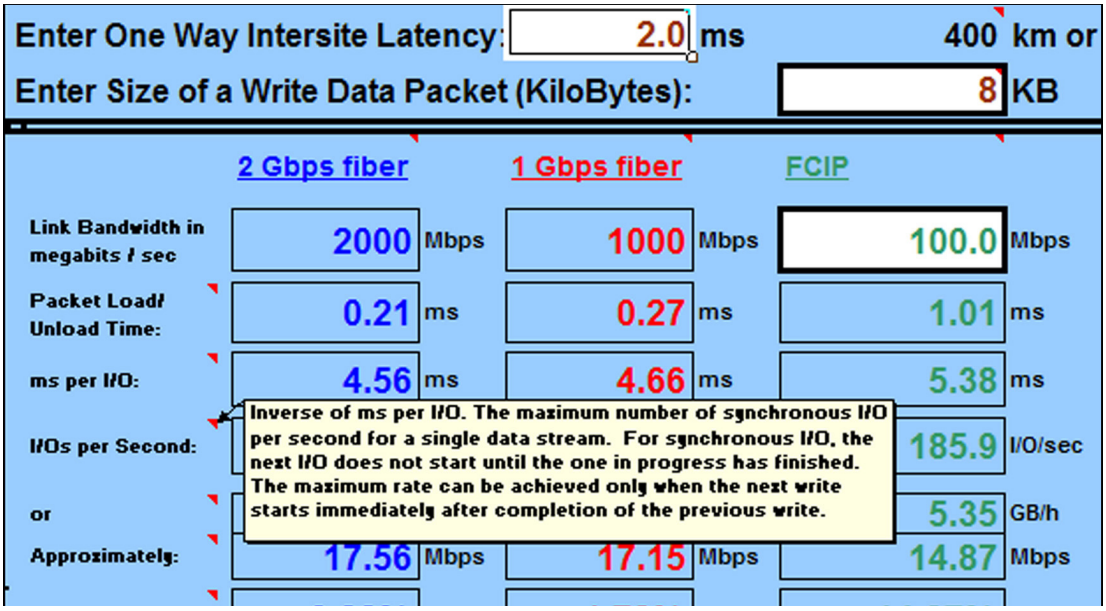

Figure 15: Output row: I/Os per second

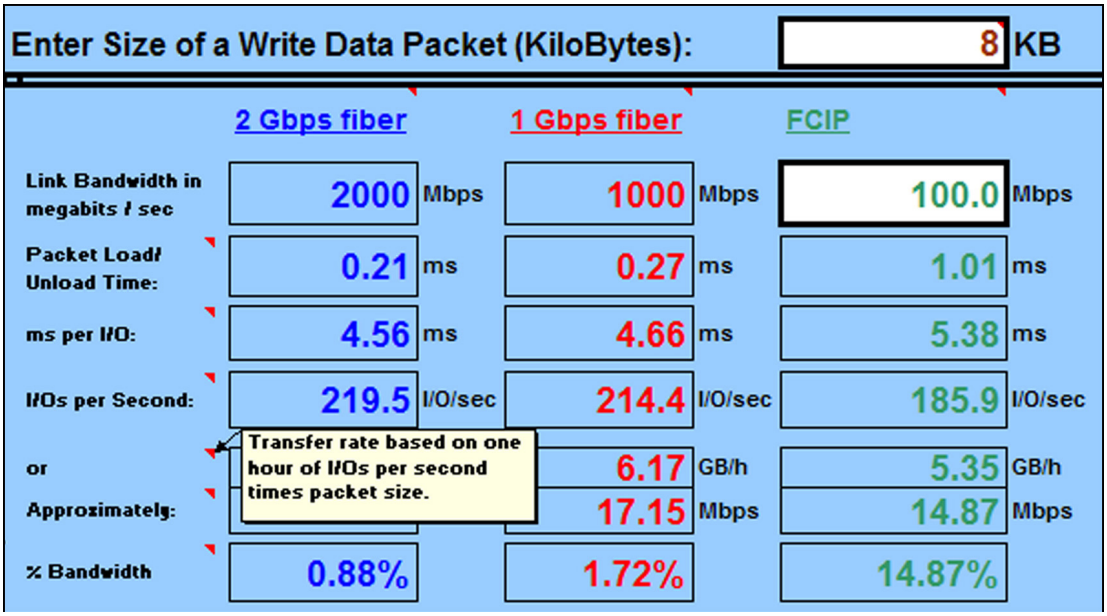

Figure 16: Output row: transfer rate in GB per hour

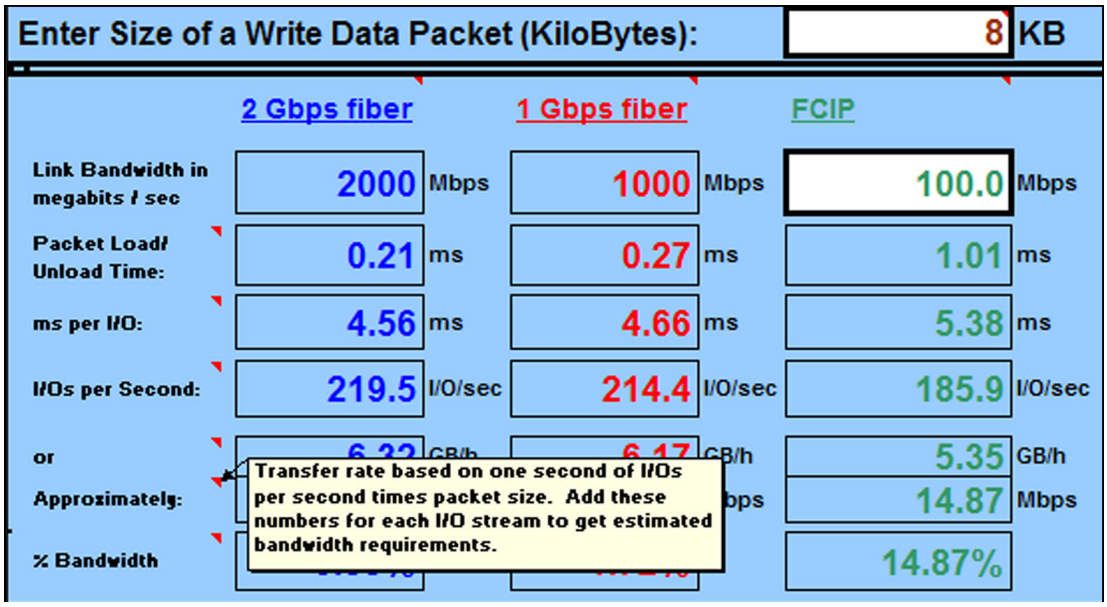

Figure 17: Output row: transfer rate in Mbps

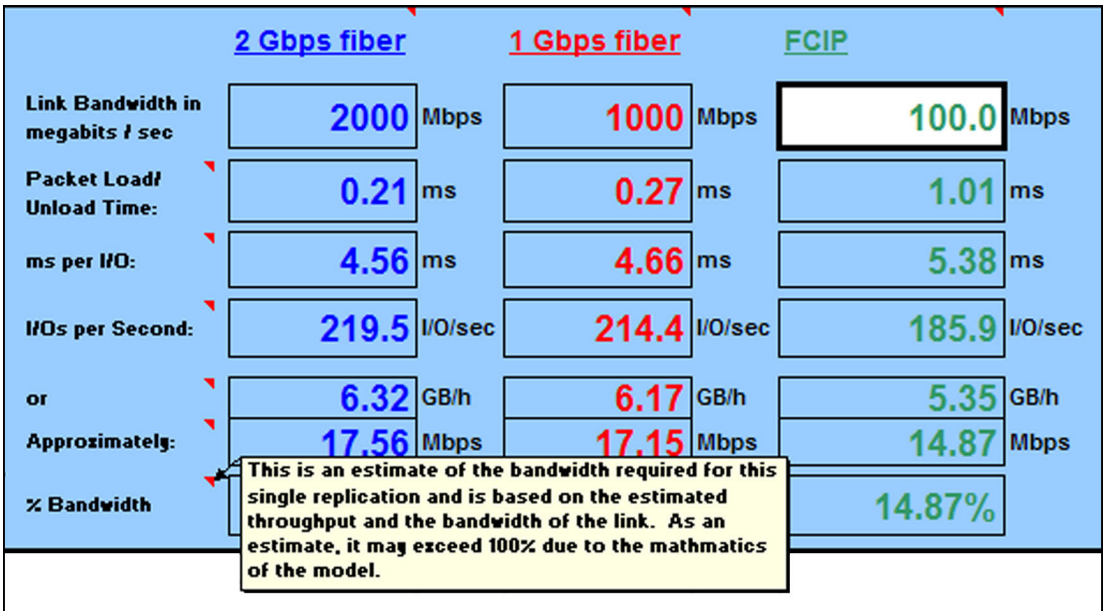

Figure 18: Output row: estimated percent bandwidth

What does this mean?

In the *HP StorageWorks Continuous Access EVA Design Reference Guide*, there are two examples, which are reproduced here. In these examples the object is to determine how much bandwidth is required to support two applications. For the examples, the proposed one-way latency is 1 ms, equivalent to 200 km of fiber optic cable. In the first example, the application uses a packet size of 32-KB (see Figure 19); in the second example, the packet size is 2-KB. Figure 19 shows that a maximum single I/O stream of 32-KB could consume approximately 50% of a 100-Mbps FCIP link. Based on best practices recommendations not to exceed 40% of a link, more bandwidth or fewer writes per second are required.

The results for the 2-KB example in Figure 20 show that even a 10 Mbps link consumes less than the recommended best practice of a maximum 40% link utilization. The data also shows that a multithreaded application issuing writes up to five times the single stream rate would be supportable using a 100 Mbps link.

Note: High data rates may be limited due to external factors such as buffer-to-buffer credits, which are beyond the scope of this tool.

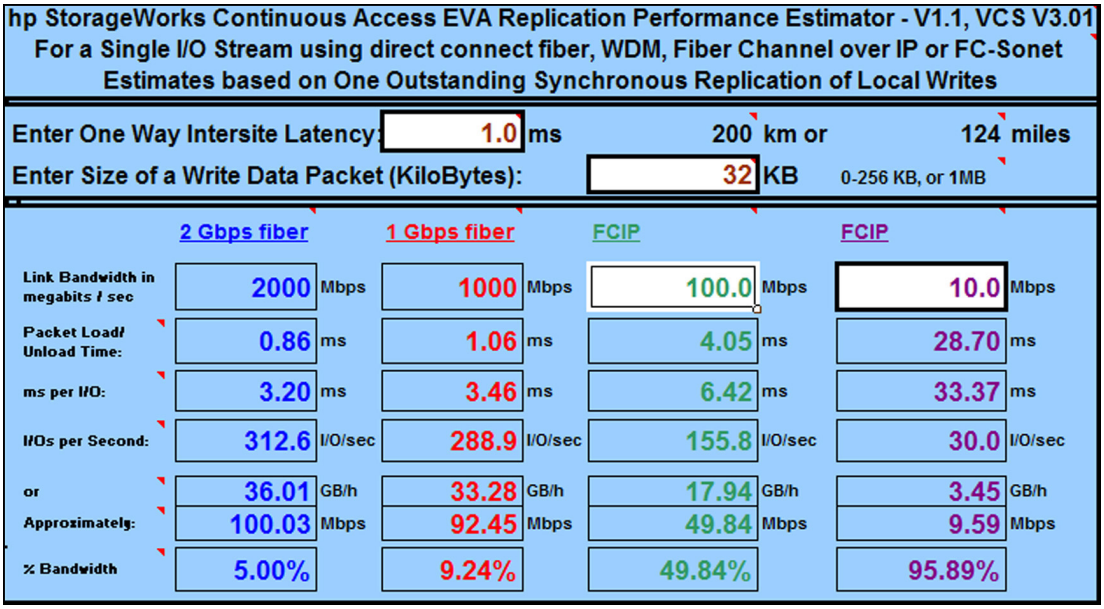

hp StorageWorks Continuous Access EVA Replication Performance Estimator - V1.1, VCS V3.01
For a Single I/O Stream using direct connect fiber, WDM, Fiber Channel over IP or FC-Sonet
Estimates based on One Outstanding Synchronous Replication of Local Writes

Enter One Way Intersite Latency: 1.0 ms 200 km or 124 miles
Enter Size of a Write Data Packet (KiloBytes): 32 KB 0-256 KB, or 1MB

	2 Gbps fiber	1 Gbps fiber	FCIP	FCIP
Link Bandwidth in megabits / sec	2000 Mbps	1000 Mbps	100.0 Mbps	10.0 Mbps
Packet Load/ Unload Time:	0.86 ms	1.06 ms	4.05 ms	28.70 ms
ms per I/O:	3.20 ms	3.46 ms	6.42 ms	33.37 ms
I/Os per Second:	312.6 I/O/sec	288.9 I/O/sec	155.8 I/O/sec	30.0 I/O/sec
or	36.01 GB/h	33.28 GB/h	17.94 GB/h	3.45 GB/h
Approximately:	100.03 Mbps	92.45 Mbps	49.84 Mbps	9.59 Mbps
% Bandwidth	5.00%	9.24%	49.84%	95.89%

Figure 19: Example with 1.0-ms latency and 32-KB packet size

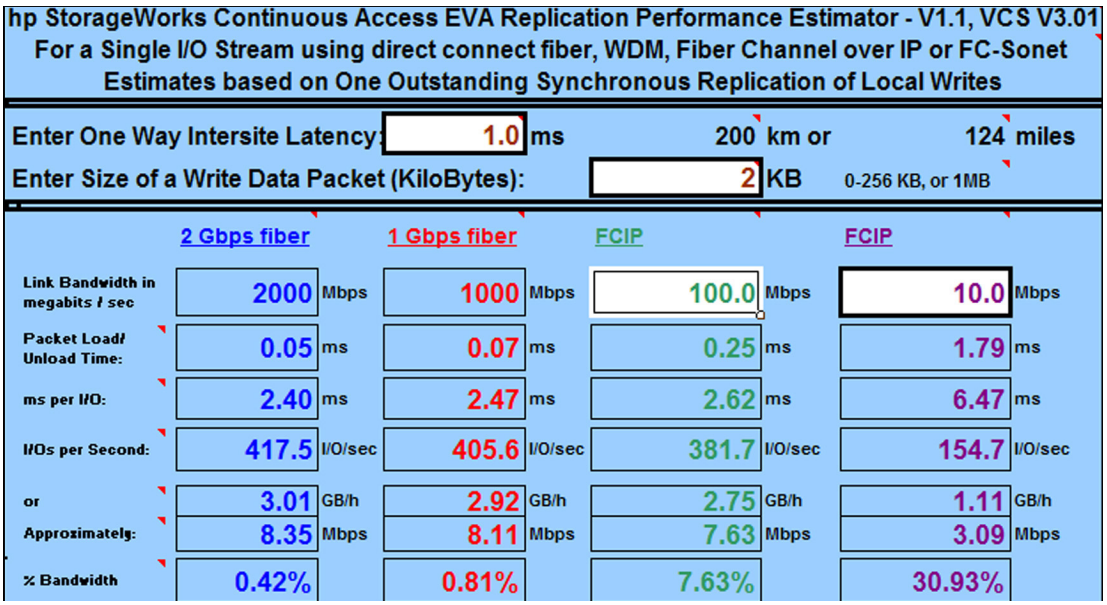

hp StorageWorks Continuous Access EVA Replication Performance Estimator - V1.1, VCS V3.01
For a Single I/O Stream using direct connect fiber, WDM, Fiber Channel over IP or FC-Sonet
Estimates based on One Outstanding Synchronous Replication of Local Writes

Enter One Way Intersite Latency: 1.0 ms 200 km or 124 miles
Enter Size of a Write Data Packet (KiloBytes): 2 KB 0-256 KB, or 1MB

	2 Gbps fiber	1 Gbps fiber	FCIP	FCIP
Link Bandwidth in megabits / sec	2000 Mbps	1000 Mbps	100.0 Mbps	10.0 Mbps
Packet Load/ Unload Time:	0.05 ms	0.07 ms	0.25 ms	1.79 ms
ms per I/O:	2.40 ms	2.47 ms	2.62 ms	6.47 ms
I/Os per Second:	417.5 I/O/sec	405.6 I/O/sec	381.7 I/O/sec	154.7 I/O/sec
or	3.01 GB/h	2.92 GB/h	2.75 GB/h	1.11 GB/h
Approximately:	8.35 Mbps	8.11 Mbps	7.63 Mbps	3.09 Mbps
% Bandwidth	0.42%	0.81%	7.63%	30.93%

Figure 20: Example with 1.0-ms latency and 2-KB packet size

Best practices

HP recommends the following best practices to avoid over-subscribing the link, a controller port, or the controller:

- Plan for utilization of one link or controller not to exceed 80% of capacity, but use other Continuous Access EVA tools to balance the load across both links and controllers.

- Plan to use no more than an average of 40% of a link when the link is load-balanced, and no more than an average of 45% at peak times.

Glossary

This glossary defines terms used in this guide or related to Continuous Access EVA and is not a comprehensive glossary of computer terms.

µm
Micrometer (10^{-6} meter).

actual disk failure protection level
The actual level of protection available to the disk group as specified in its current configuration.

allocation policy
Storage system rules that govern how virtual disks are created. There are two rules:

- Allocate Completely—The space a virtual disk requires on the physical disks is reserved, even if the virtual disk is not currently using the space.
- Allocate on Demand—The space a virtual disk requires on the physical disks is not reserved until needed.

alternate site
See remote site.

asynchronous replication
See DR group write mode.

asynchronous transfer mode (ATM)
Communications networking technology for LANs and WANs that carries information in fixed-size cells of 53 bytes (5 protocol and 48 data).

bandwidth
The transmission capacity of a link or system, usually measured in bits per second.

bidirectional
The circumstance when a storage system is configured so that it contains both source and destination Vdisks.

Bit Error Rate (BER)
The rate at which a single bit error is detected in communications networks.

CAC
Corrective action code. A display component that defines the action required to correct a problem. This component is displayed on the HP StorageWorks Command View EVA graphical user interface (GUI).

CA
Continuous Access. A storage-based HP StorageWorks solution consisting of two storage systems performing storage array-to-storage array replication, along with the management user interface (Continuous Access UI) that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.

CGR
See constant bit rate.

clone
A clone is a physical, block-for-block copy of an entire volume. The cloning capabilities reside in the HSV controller. The cloning operation creates a third copy of the parent volume. After the clone is created, it acts as a local mirror until data on the clone has been normalized and the clone broken off. Once broken off, the clone can be mounted on another host that is connected to the controller.

Command View EVA
The Command View EVA consists of:

- The graphical user interface (GUI) that displays the usage of the storage system.
- The software behind the GUI, which controls and monitors the functions.

The Command View software can be installed on more than one HP OpenView Storage Management Appliance in a fabric. Each installation of the Command View software is a management agent. The client for the agent is a standard browser.

console LUN
A SCSI-3 virtual object that makes a storage array accessible by the host before any virtual disks are created. Also called a *communication LUN*.

console LUN ID
The ID that can be assigned when a host operating system requires a unique ID. The console LUN ID is assigned by the user, usually when the storage system is initialized. *See also* console LUN.

constant bit rate (CBR)
Category of ATM Quality of Service (QoS) that supports a constant or guaranteed data rate. CBR supports applications that need a highly predictable transmission rate.

Continuous Access
See CA.

Continuous Access UI
A tool for managing the replication of storage objects in a SAN. It provides a graphical user interface for the management of disk I/O, failover, and maintenance operations.

controller heartbeat LED
A green LED that flashes on the controller OCP to indicate that the HSV controller is operational.

copy set
Generic term for a logical disk in one storage array that is replicated to another logical disk in another storage array. There are two states: *normal* and *copying*. The term is commonly used to represent the pair of Vdisks, one on the source array and one on the destination array.

corrective action code
See CAC.

data distribution
Pushing copies of data to geographic locations to make it more easily accessible to many customers.

data entry mode
The state in which controller information can be displayed or controller configuration data can be entered. On the Enterprise storage system, the controller data entry mode is active when the LCD on the HSV controller OCP is flashing.

data migration
Moving data to a new location or to a logical disk with a different capacity.

data mining
A process that makes data available so that undiscovered and useful information can be extracted, analyzed, or tested.

data movement
Activities such as data backup, data migration, and data distribution.

data replication mode
See DR (data replication) mode.

default disk group
The first disk group created when the storage system is initialized. The default disk group can contain up to the entire set of physical disks in the array. The minimum number of physical disks the default disk group can contain is eight. The maximum is the number of installed disks and is limited to 14 drives per attached drive enclosure.

destination Vdisk
A Vdisk that is the recipient of replicated data from a source Vdisk.

destination Vdisk access
A storage system's ability to allow host access to a destination Vdisk. There are three options:

- Disabled
- Read-Only (write-protected; for future use)
- Active/Active (for future use)

disaster recovery (DR)
The ability to respond to an interruption in services by implementing an in-place plan for restoring an organization’s critical business functions. Note that DR is used in context to mean either disaster recovery or data replication.

disaster tolerance (DT)
The capability for rapid recovery of user data from a remote location when a significant event or disaster occurs at the local computing site. It is a special combination of high-availability technology and services that can continue the operation of critical applications in the event of a site disaster. DT systems are designed to allow applications to continue operating during the disaster recovery period.

disk failure protection
For each disk group, the controllers hold in reserve the space in the physical disk pool equivalent to a selected number of physical disk drives.

The three levels of disk failure protection are:

- None—No protection is present.
- Single— The capacity of one physical disk is reserved.
- Double—The capacity of two physical disks is reserved.

Disk failure protection occurs when the storage array sets aside reserved capacity to take over the functionality of a failed or failing physical disk drive. In groups with mixed capacity disk drives, the reserved capacity is based on the largest disk in the disk group. The system must cover a failure in any drive, so it reserves enough capacity to cover the largest failure that could happen.

disk group
A named group of physical disks selected from all the available physical disks in the storage system where one or more Vdisks can be created. A physical disk can belong to only one disk group.

disk group occupancy alarm level
An event code that is generated when the amount of data stored in the disk group reaches a peak level of the total disk group capacity. For example, if the disk group capacity is 288 GB, and the occupancy alarm level is 80%, the event code is generated when the amount of data in the disk group reaches 230.4 GB. The default occupancy alarm level is 95% (274 GB in this example) of the total disk group capacity.

disk migration state
The state of data on the physical disk drive. Two states are possible:

- Stable—The physical disk drive has not failed.
- Migrating—The disk drive is failing or failure is predicted. If this state occurs, data is moved by the controller from the failing disk to other disk drives in the same disk group.

disk replacement delay
The time that elapses between detection of a possible drive failure and when the controller starts searching for spare disk space. Drive replacement seldom starts immediately, in case the failure was a glitch or temporary condition.

DR (data replication) group
A VCS construct organizing one or more Vdisks in an HSV storage system so that they replicate to the same specified destination, failover together if a single Vdisk in the collection fails, and preserve write ordering within the collection. This is the HSV implementation of an association set, and its member Vdisks are the HSV implementation of copy sets. Note that DR is used in context to mean either disaster recovery or data replication.

DR (data replication) group direction
The replication direction of a DR group. There are two states:

- Original (from *Home*)
- Reversed (failed over, or toward *Home*)

DR (data replication) group log state
The current behavior of the log associated with a DR group. In the state options, references to multiple destinations are for future use. There are three possible states:

- Idle—No destination is logging or merging.
- Logging—At least one destination is logging; none are merging.
- Merging—At least one destination is merging.

DR (data replication) group write mode
Characterizes how a write from a host is replicated. A DR group has two modes of replicating writes from the source Vdisk to the destination Vdisk:

- Synchronous Replication—A host write is acknowledged as complete by the storage array after the write command and data are inserted into the source and destination storage array cache. If the destination array is not available or the DR group is merging, the write command and data is inserted into the tail of the write history log.
- Asynchronous Replication—(For future use) A host write is acknowledged as complete by the storage array after the write command is inserted into the source array cache.

DR (data replication) mode
The operational mode of a DR group that indicates the capability of I/O to be written to its source or its destination, or to both. There are four options:

- Active/Passive (source)
- Active/Passive (destination)
- Active/Active (master)
- Active/Active (non-master)

Both Active/Active modes are for future use.

dual fabric
Two independent fabrics providing multipath connections between FC end devices.

E1
The standard European carrier for transmission at 2.048 Mbps.

E3
The standard European carrier for transmission at 34.368 Mbps.

E4
The standard European carrier for transmission at 139.264 Mbps.

event
Any change that is significant to the storage system. Events include:

- A state change in hardware
- A state change in a logical element, such as a virtual disk
- A completion of a procedure
- An environmental change
- An operational failure
- A change in configuration, such as a new virtual disk that has been created or a new physical disk has been inserted

fabric
A network of at least one Fibre Channel switch and attached devices.

failover
This term is context specific.

- **DR (data replication) group failover—**An operation to reverse the direction of a DR group.
- **Managed set failover—**An operation to reverse the direction of all the DR groups in the managed set.
- **Fabric or path failover—**The act of transferring I/O operations from one fabric or path to another.
- **Controller failover—**When a controller assumes the workload of its partner.

failsafe mode
When enabled, this is a data replication group mode in which all source Vdisks in the group become both unreadable and unwritable if any of their corresponding member Vdisks is unreachable. No write history logging takes place for the Vdisks. There are two states: *enabled* and *disabled*.

failsafe-locked
A condition that prohibits any further host I/O to any of the Vdisks in the data replication (DR) group. A failsafe-enabled DR group becomes failsafe-locked whenever it is unable to complete a write to its destination, or when a member Vdisk fails. It requires immediate intervention.

FC connection
A Fibre Channel path between two storage systems or a host and its storage. A connection is made up of multiple FC links.

FC link
Fibre Channel Link. A path between two adjacent FC ports.

FCIP (Fibre Channel over Internet Protocol)
A method of moving Fibre Channel data over a standard Internet Protocol network. The IP network may be Ethernet, SONET, or ATM based.

Fibre Channel (FC)
Technology for very high speed, switching-based serial transmissions.

Fibre Channel connection
See FC connection.

flush cache
The act of writing data from cache to storage media.

full copy
A copy operation for which all 1-MB blocks written on a source Vdisk since it was created are replicated to a destination Vdisk.

Gbps or Gb/sec
Gigabits per second. A measurement of the rate at which the transfer of bits of data occurs. Nominally, 1-Gbps is a transfer rate of 1,000,000,000 (10^9) bits per second.

For Fibre Channel transceivers or FC loops the Gb transfer rates are:

- 1-Gbps is a transmission rate of 1,062,500,000 bits per second.
- 2-Gbps is a transmission rate of 2,125,000,000 bits per second.

gigabit interface converter (GBIC)
The hardware devices inserted into the ports of the Fibre Channel switch that hold the fiber optic cables in 1-Gb fabrics. GBIC devices are available for short-range applications (0.5 to 500 m), long-range applications (up to 10 km), and very long distances (up to 100 km).

gigabit link module (GLM)
Permanently installed GLMs provide fiber optic cable transmission at distances of 0.5 to 500 m, depending on the cable size and quality of its installation.

heterogeneous SAN support
The ability for the product to operate with different operating systems and storage systems in a SAN.

hierarchical storage virtualization (HSV)
A transparent abstraction of storage at the block level. Using block-level mapping techniques, storage virtualization presents host servers with a logical view of storage in the form of virtual disks. All raw storage is pooled, and virtual disks (or Vdisks) are created that draw their capacity from the pool.

high availability (HA)
Redundant systems, software, and IT processes to reduce the risk of downtime. No single points of failure.

high-availability switches
Fibre Channel Director and Edge switches. See the *HP StorageWorks Continuous Access EVA Design Reference Guide* for a list of Continuous Access EVA-supported devices of this type.

home
The preferred storage system for the source Vdisks of a data replication group. By default, home is the storage system on which a source is created, although this designation is user settable.

homogeneous SAN support
Implies the ability for the product to operate with homogeneous operating systems and homogeneous storage systems in a SAN.

hop
One interswitch link.

HSV
See hierarchical storage virtualization.

I/O module
Input/output module. The enclosure element that is the FC-AL interface to the host or controller. I/O modules are bus speed specific: either 1-Gbps or 2-Gbps.

IDX
A 2-digit decimal number portion of the HSV controller termination code display that defines one of 32 locations in the termination code array that contains information about a specific event.

in-band communication
Communication that uses the same communications pipe as the operational data. *See also* out-of-band communication.

infrastructure switches
Fibre Channel core and SAN switches. See the *HP StorageWorks Continuous Access EVA Design Reference Guide* for a list of Continuous Access EVA-supported devices of this type.

initialization
A configuration step that binds the controllers together as a storage array and establishes preliminary data structures on the disk array. Initialization also sets up the first disk group, called the default disk group, and makes the storage system ready for use.

IOPS
I/O operations per second.

ISL
Intersite link or interswitch link. The abbreviation is context sensitive.

license key
A license key is required to operate the HSV controller software. The license key is a WWN-encoded sequence that is obtained from the HP license key fulfillment website. Three types of license keys exist:

- Basic—This license key is required to unlock VCS. The basic license key covers both controllers in the storage system. The basic license key is required for system initialization.
- Snapshot— This license key is needed to unlock the snapshot and Snapclone features. This license can be added any time *after* the system has been initialized.
- Data Replication Management—This license is necessary to activate the data replication features of the HSV controller software. It is also necessary for running the Continuous Access UI.

link
A connection between two adjacent Fibre Channel ports, consisting of a pair of simplex fibers—one for each direction of information flow. An example is the connection between the Fibre Channel switch port and the HSV controller.

local site
One or more storage arrays residing at the local data center, which contains the source data. A storage system may be considered a local site for some Vdisks and a remote site for other Vdisks.

log
Storage that is used for logging.

logging
Usually context sensitive. In Continuous Access EVA, *logging* refers to the history of host write commands (and data) when the destination array is not accessible. If failsafe mode is enabled, logging does not occur.

logical unit number
See LUN.

longwave GBIC
A GBIC used with 9-µm single mode fiber and operating on a wavelength of 1350 nm. The 1-Gbps model has SC-style connectors and runs to a distance of 10 km. The 2-Gbps SFP has LC-style connectors and runs to a distance of 30 to 40 km. Also known as an *LD* (long-distance) *GBIC*.

LUN
Logical unit number. An identifier through which a Vdisk is presented to a host.

managed set
Any collection of data replication (DR) groups selected by the user for the purpose of managing them. For example, a managed set can be created to manage all DR groups whose sources reside in the same cabinet or all DR groups that deal with a particular set of applications.

master
The controller in the storage array that powers up first. Also called *storage system master*.

Mbps
Megabits per second. Typically, a unit of measurement for bandwidth of a serial link.

MBps
Megabytes per second. Typically, a unit of measurement for throughput for a serial link.

merge
See merging.

merging
Transferring the contents of the write history log to the destination Vdisk to synchronize the source and destination Vdisks.

MMF
Multimode fiber, typically 50-µm, although 62.5-µm is also supported at reduced distances.

µm
Micrometer (10^{-6} meter).

nm
Nanometer (10^{-9} meter).

OC1
The optical carrier signaling rate synchronous data handling (SDH) (STM0)/SONET of 51.84 Mbps.

OC3
The optical carrier signaling rate synchronous data handling (SDH) (STM0)/SONET of 155.52 Mbps.

OC12
The optical carrier signaling rate synchronous data handling (SDH) (STM0)/SONET of 622.08 Mbps.

OC48
The optical carrier signaling rate synchronous data handling (SDH) (STM0)/SONET of 2488.32 Mbps.

OCP
Operator control panel. The element on the front of an HSV controller that displays the controller's status using LEDs and an LCD. Information selection and data entry are controlled by the OCP push buttons.

operation state
Current operating condition of a system component. There are three states:

- Normal
- Failed
- Attention (indicates possible problem)

out-of-band communication
Communication that uses a different communications pipe than that used by operational data.

See also in-band communication.

peak cell rate (PCR)
Peak cell rate is the maximum transmission speed of a virtual connection and is a required parameter for the CBR Quality of Service (QoS) category.

permanent virtual circuit (PVC)
Logical connection between two points that are manually defined by the network administrator.

preferred path
A preference for which controller of the storage array manages the virtual disk. This preference is set by the user through Command View EVA when creating the virtual disk. The primary purpose of preferring a path is load balancing.

presentation
The process whereby a controller presents a virtual disk only to the host computer that has authorized access.

primary site
See local site.

quality of service (QoS)
Each virtual connection in a communications network has a service category. The performance of the connection is measured by the established QoS parameter.

reconstruction
The process of regenerating the contents of a failed member data. The reconstruction process writes the data to a spare set disk and incorporates the spare set disk into the mirrorset, striped mirrorset, or RAIDset from which the failed member came.

relationship
The arrangement created when two storage systems are partnered for the purpose of replicating data between them.

remote site
A site containing one or more storage arrays that contains copies of data stored at a local site. A remote site could be in the same room or building as the local site, but usually it is not.

reservation state
Three possible reservation states exist for a virtual disk:

- None—No host has a reservation on the virtual disk.
- Regular—One host has a reservation on the virtual disk. A regular reservation will not be preserved through failovers or power failures.
- Persistent—One or more hosts have the virtual disk reserved. A persistent reservation is normally preserved through failovers and power failures. A persistent reservation can be released by other hosts.

resume
Command issued to a data replication (DR) group or managed set that causes replication to resume after being suspended. This command may initiate a merging of the DR group log or a full copy. *See also* suspend.

SCSI (small computer system interface)
An American National Standards Institute (ANSI) interface standard defining the physical and electrical parameters of a parallel I/O bus. A processor-independent standard protocol for system-level interfacing between a computer and intelligent devices including hard drives, disks, CD-ROM drives, printers, scanners, and other devices.

SFP (small form factor pluggable GBIC)
A 2-Gbps GBIC. Typically uses an LC connection.

shortwave GBIC
A GBIC used with 50- or 62.5-µm multimode fiber and operating on a wavelength of 850 nm. The 1-Gbps model has SC-style connectors and runs to a maximum of 500 m. The 2-Gbps SFP has LC-style connectors and runs to a maximum of 300 m. Also known as an *SD* (short distance) *GBIC*.

single path
A single connection or path between storage systems containing source and replicated data, or between a host and the storage assigned to that host.

site failover
Command to change a destination role to a source role at the designated site. *See also* failover.

slave
The controller that powers up last. Also called *storage system slave*.

SMF
Single-mode fiber. A type of cable, typically 9-µm, although 8-µm and 10-µm are also supported.

Snapclone
A snapshot or virtual copy of the data, which becomes a clone or physical copy of the source over time.

snapshot
A temporary Vdisk that reflects the contents of another virtual disk at a particular point in time. A snapshot operation is performed only on an active virtual disk. The active disk and its snapshot constitute a virtual family.

SONET (Synchronous Optical Network)
An ANSI standard for transmitting bits over fiber optic cable. See OC1, OC3, OC12, and OC48.

source Vdisk
Contains original data that is replicated to its destination.

Storage Management Appliance
A single aggregation point for data management. The HP OpenView Storage Management Appliance is a specialized server on which SAN applications, the Element Manager, and the Continuous Access UI run.

storage pool
The aggregated blocks of available storage in the total physical disk array.

storage system
An array of HSV controllers and the array of physical disks they control. A storage system may contain Vdisks that are sources as well as Vdisks that are destinations. Sometimes used in context as follows:

- Source storage system—Used in the context of a particular data replication (DR) group, this is the storage system in which the source Vdisk resides.
- Destination storage system—Used in the context of a particular DR group, this is the storage system at which the destination Vdisk resides.

storage virtualization
The transparent abstraction of storage at the block level. It separates out logical data access from physical per disk, or per array, data access.

suspend
Command issued to a data replication (DR) group or managed set that temporarily halts replication of I/O from all source Vdisks to destination Vdisks in that DR group. Source Vdisks continue to run I/O locally, and the I/O is also copied to the DR group log. May not be issued if the DR group is failsafe enabled. *See also* resume.

synchronous mode
The mode of operation of a data replication (DR) group where the data is written to the source and destination caches, after which a completion acknowledgement is sent to the host. *See also* DR group write mode.

synchronous replication
See DR group write mode.

T1
The standard North American carrier for transmission at 1.544 Mbps.

T3
The standard North American carrier for transmission at 44.736 Mbps.

throughput
In data transmission, a measure of performance for the amount of data moved successfully from one place to another in a given time.

trunking
The combining of two or more low-speed links into one virtual high-speed link. In a Fibre Channel fabric, trunking means combining two or more intersite links (ISLs) into one virtual high-speed ISL.

unspecified bit rate (UBR)
Offers no traffic-related Quality of Service (QoS) guarantees. Not acceptable for FCIP use.

UUID
Universal unique identifier. A unique 128-bit identifier associated with HSV objects.

VCS
Virtual Controller Software. The firmware that runs the storage system.

Vdisk
Virtual disk accessible from hosts attached to the SAN. When it is a member of a data replication group, a Vdisk is the HSV implementation of a copy set, and it can have two states: *normal* and *copying*.

very long-distance GBIC
A GBIC used with 9-µm single-mode fiber and operating on a wavelength of 1500 nm. The 1-Gbps model has SC-style connectors and runs to a distance of 100 km (160 km with two hops). The 2-Gbps model has LC-style connectors, but is not available at the time of publication. Also known as a *VLD GBIC*.

virtual channel (VC)
The lowest-order logical address in asynchronous transfer mode of an ATM network. VC refers to a given circuit on a link.

virtual channel identifier (VCI)
The field of the cell header that stores the VC address.

Virtual Controller Software
See VCS.

virtual disk
See Vdisk.

virtual path (VP)
The highest-order part of a logical address in ATM networks. VP refers to a given group of circuits on a link.

virtual path identifier (VPI)
The field of the cell header that stores the VP address.

Vraid0
A virtualization technique that provides no data protection. The data host is divided into chunks and distributed on the disks that constitute the disk group from which the virtual disk was created. Reading and writing to a Vraid0 virtual disk is very fast and makes the fullest use of the available storage but provides no data protection (redundancy) unless there is parity.

Vraid1
A virtualization technique that provides the highest level of data protection. All data blocks are mirrored, or written twice, on separate physical disks. For read requests, the block can be read from either disk, which can increase performance. Mirroring requires the most storage space because twice the storage capacity must be allocated for a given amount of data.

Vraid5
A virtualization technique that uses parity striping to provide moderate data protection. Parity is a data protection mechanism for a striped virtual disk, one on which the data to and from the host is broken down into chunks and distributed to the physical disks comprising the disk group in which the virtual disk was created. If the striped virtual disk has parity, another chunk (a parity chunk) is calculated from the set of data chunks and written to the physical disks. If one of the data chunks becomes corrupted, the data can be reconstructed from the parity chunk and the remaining data chunks.

wavelength division multiplexing (WDM)
The ability to have multiple optical signals share a single optical cable.

world wide LUN name (WWLN)
The 128-bit identifier associated with a Vdisk (64 bits come from the controller WWN).

world wide name (WWN)
A 64- or 128-bit identifier used to uniquely identify the address of a component on the fabric.

WWLN
See world wide LUN name.

WWN
See world wide name.